Improving Risk Management

RIESGOS-CM

Análisis, Gestión y Aplicaciones

P2009/ESP-1685

Technical Report 2010.17

Sensitivity Analysis

Andres Redchuk and David Ríos Insua

http://www.analisisderiesgos.org

Sensitivity Analysis

Andres Redchuk, Dept.of OR and Stats, Rey Juan Carlos University
David Ríos Insua, Royal Academy of Sciences, Spain

1.   Introduction

In Operations Research, sensitivity analysis describes the methods and tools used to study how the output of a model varies with changes in the input data. The input data may refer to parameters affecting the objective functions and/or constraints or to the structure of the problem. Depending on the problem and model, the output could refer to:

- the optimal alternative and/or the optimal value, or,

- a set of alternatives with a certain property. Some examples include the non-dominated set in a multi-objective optimization problem; the set of alternatives satisfying certain constraints in a classification problem; or the set of the, say, five best alternatives.

Typical questions addressed within sensitivity analysis are whether a given optimal solution will remain as such if inputs are changed in a certain way, and, if not, which other alternatives may become optimal. Finding the most critical directions for changes in inputs that may affect the model output are also relevant sensitivity analysis issues, see French and Ríos Insua (2000), Saltelli et al (2000), and Saltelli et al. (2004) for reviews.

As motivating examples, consider these:

1. Consider a linear programming problem. One may be interested in checking how the costs (reduced costs) and/or the right hand side terms and/or the technological matrix terms impact over the optimal solution. A typical question would be: does the optimal solution change if one of the costs increases so much?

2. Consider a decision analysis problem. One may be interested in checking the impact of the beliefs and preferences, modeled, respectively, through a probability distribution and a utility function, over the optimal alternative and its expected utility. For example, one could wonder how much a relevant binomial parameter and a risk tolerance parameter should be changed to make a certain alternative optimal.

3. Consider now a multi-objective decision making problem solved with a multi-attribute value function. One may want to check which alternatives

become optimal when the weights vary within a range around the current weight settings.

An important, and occasionally controversial, issue in sensitivity analysis is the distinction between decision sensitivity and value sensitivity, see Kadane and Srinivasan (1996). A variety of situations may hold. For instance, when performing sensitivity analysis, it may happen that value changes considerably with virtually no change in the optimal alternative.

2. Motivations

There may be many reasons to check the sensitivity of the output of an OR model to its inputs. A first reason may be the almost ubiquitous uncertainty in the inputs. We may not be willing or capable of assessing such uncertainty with a probability distribution. Then, baseline values for the inputs could be assessed and changes in how they affect the output are observed.

Similarly, the assessment of the inputs might be affected by inherent imprecision and output robustness may be checked. In relation with this, it may be interesting to check the robustness of the output under various input scenarios.

Note also that, since some of the inputs to an analysis may encode the decision makers' (DM) subjective judgments, their implications and possible inconsistencies should be explored. The need for sensitivity analysis is further emphasized by the fact that the assessment of such judgments could be a difficult task. For example, it is frequently mentioned that assessing a subjective probability distribution is involved. Consider the simplest case in which it is desired to elicit a prior over a finite set of states $\theta_i, i \in \{1, \dots, I\}$. A common technique to assess a precise probability distribution $\pi(\theta_i) = p_i$ proceeds as follows, with the aid of a reference experiment: one progressively bounds $\pi(\theta_i)$ above and below until no further discrimination is possible and, then, takes the midpoint of the resulting interval as the value of p_i. Instead, one could directly operate with the obtained constraints $\alpha_i \leq \pi(\theta_i) \leq \beta_i$, acknowledging cognitive limitations. This is an especially important point, as the DM's judgments will evolve through the analysis until they are requisite. Sensitivity analysis may guide such process.

In relation with the limitations of elicitation, consider also the situation in which there are several decision makers and/or experts involved in the elicitation. Then it is not even necessarily possible theoretically to obtain a single model: one might be left with only classes of each, corresponding to differing expert opinions, and we may need to study the model under those various settings.

Finally, note that sensitivity analysis may be used to perform value of information calculations that allow us to compute how much should we pay for information used to reduce our uncertainty in an analysis.

To sum up, sensitivity analysis aims at increasing the confidence in an OR model and its output,by providing an understanding of the responses of the model to changes in the inputs.

3. Foundations

A number of results show that we may deal with imprecision in model inputs through a class of probability distributions and a class of utility functions. These results have two basic implications. First, they provide a qualitative framework for sensitivity analysis, describing under what conditions the standard and natural sensitivity analysis approach of perturbing the initial input assessments within some reasonable constraints may be undertaken. Second, they point out to the basic solution concept of robust approaches, thus indicating a key computational objective in sensitivity analysis, as long as the interest is in decision analytic problems: that of *non-dominated alternatives*. An alternative a dominates another alternative b, if its evaluation is better for each potential input to the analysis. Then, an alternative a is *non-dominated* if no other feasible alternative dominates it. This corresponds to a Pareto ordering of alternatives based on inequalities on their evaluations.

To construct an appropriate framework for general sensitivity analysis, the standard decision theoretic axiomatic foundations should be reconsidered to account for imprecision in model inputs. Although this approach does not lead, so far, to such a well rounded development as in the precise case, in which various axiomatizations essentially lead to the subjective expected utility model, various partial results do exist, see e.g., Ríos Insua (1990) or Walley (1991), leading to, essentially, the same conclusion: imprecise beliefs and preferences may be modeled by a class of priors and a class of utility functions, so that preferences among alternatives may be represented by inequalities of the corresponding posterior expected utilities. The basic argument for such results assumes that the underlying preference relation, rather than being a weak order (complete and transitive), is a quasi order (reflexive and transitive9.

4. Key approaches to sensitivity analysis

Clearly, as there is a large variety of OR models, there is a comparatively large number of approaches to sensitivity analysis. Only a few of the approaches that may be applied to various OR models are described, without much dwelling into their numerical details.

4.1. Testing alternative inputs

One first generic approach refers to changing the inputs to the model and observe variations in the output. This may be done in several ways.

4.1.1. Trying other values

The first approach may be termed the *informal* one, which considers several inputs and compares the quantity of interest (e.g., the posterior mean, the difference in value between two alternatives, the optimal alternative) under them. The approach is very popular because of its simplicity. While this is a healthy

practice and a good way to start a sensitivity analysis, in general this will not be sufficient and more formal analysis should be undertaken: the limited number of priors chosen might not include some which are compatible with the prior knowledge and could lead to very different values of the quantity. Sometimes, the alternative inputs considered are randomly generated.

4.1.2. Parametric analysis

Another way of changing the inputs is through parametric analysis. Using a baseline input assessment, one determines a relevant direction to perturb the inputs and considers a parametric perturbation along such direction observing whether there is a change, or not, in, e.g., the optimal alternative, see Gal and Greenberg (1997).

4.1.3. Global robustness

Another popular approach in SA is called *global sensitivity*. All inputs compatible with the prior knowledge available are considered and robustness measures are computed as the inputs vary within that class. Computations are not always easy since they require the evaluation of suprema and infima of quantities of interest. The choice of the class of inputs should be driven by the following goals:

1. the class should be related with the elicitation method used;

2. the class should contain only "reasonable" inputs, avoiding unreasonable inputs which might erroneously lead to lack of robustness;

3. computation of sensitivity measures should be as simple as possible.

The robustness measures provides, in general, a number that should be interpreted in the following way:

- if the measure is "small", then robustness is achieved and any input in the class can be chosen without relevant effects on the quantity of interest;

- if the measure is "large", then new data should be acquired and/or further elicitation to narrow the class, recomputing the robustness measure and stopping as before;

- otherwise, if the measure is "large" and the class cannot be modified, then an input can be chosen in the class but the relevant influence of our choice over the quantity of interest should be considered carefully.

Given a class of inputs, global sensitivity analysis will usually pay attention to the range of variation of a quantity of interest as the input ranges over the class. As an example, in a decision theoretic problem, suppose a quadratic loss function is used in a problem. The optimal rule is the posterior expectation. If there is imprecision about the prior, the range of the posterior expectation as the prior ranges in the class would be computed.

4.2. Behavior of the optimal alternative

The other family of sensitivity analysis approaches studies the behavior of the output of interest under small input perturbations, either via differential approaches or convergence arguments.

4.2.1. Local sensitivity

Local sensitivity analysis studies the rate of change in inferences and decisions, using functional analysis differential techniques, trying to assess how a small change in the input affects the quantity of interest. The two issues involved refer to choosing the derivative and the corresponding norm, over the appropriate class of inputs. For the first choice, Fréchet derivatives, total derivatives and Gateaux differentials have been used, among others. Divergence measures have been used as well. For the second choice, the total variation, Prohorov, Levy, and Kolmogorov metrics have been used among others. The direction providing the supremum norm is used as the most sensitive direction; alternatively, the average sensitivity is sometimes used integrating the norm along all possible relevant directions.

As an example, Ruggeri and Wasserman (1993) measured the local sensitivity of a posterior expectation with respect to the prior by computing the norm of the Fréchet derivative of the posterior with respect to the prior over several different classes of inputs.

4.2.2. Stability

Stability theory provides another unifying, general sensitivity framework, formalizing the idea that imprecision in elicitation of inputs should not affect the optimal decision greatly. When *strong stability* holds, a careful enough elicitation leads to decisions with optimal value close to the greatest achievable; when *weak stability* holds, at least one stabilized decision will have such property. However, when neither concept of stability applies, even small elicitation errors may lead to disastrous results in terms of large losses in value.

Stability theory studies the convergence of decisions, nearly optimal for input sequences converging to the baseline inputs, to the corresponding optimal alternative. The arguments involved refer to the continuity of the relevant operator, e.g. the posterior expected utility functional. Note that stability is not always guaranteed, even in standard problems, as shown e.g. in Kadane *et al* (1996). Such counterexamples show a need for conditions which ensure stability. While these conditions simplify the task of verifying stability, it can still be hard to do so in practice.

4.3. An operational approach to sensitivity analysis

An operational approach to sensitivity analysis in OR models may be described as follows. At a given stage of the analysis, information on the DM's inputs is elicited, and the class of all inputs compatible with such information

is considered. The set of non-dominated solutions is approximated. If these alternatives do not differ too much in their value, the analysis may be stopped. Otherwise, additional information will be gathered. This would further constrain the class: the set of non-dominated alternatives will be smaller. It is hoped that this iterative process would converge until the non-dominated set is small enough to reach a final decision. It is conceivable in this context that at some stage it might not be possible to gather additional information yet there remain several non-dominated alternatives with very different values. In these situations, ad hoc approaches such as maximin solutions may aid as a way of selecting a single robust solution: each alternative is associated with its worst evaluation, given the current input imprecision. The alternative with best worst evaluation is suggested. Alternatively, a prior over the class of inputs could be built and base choice on expectations over evaluations.

The relevant steps are now described.

4.3.1. Non-dominated alternatives

As mentioned, a key solution concept is the efficient set, i.e., the set of non-dominated alternatives. In most cases, it is not possible to compute the non-dominated set exactly, and thus approximation schemes are necessary. Typically, one would proceed by randomly sampling the set of alternatives, randomly sampling the set of inputs, and checking dominance among pairs of alternatives. Under appropriate conditions, this sampling scheme is such that the sample non-dominated set converges to the non-dominated set, as the sample size grows.

4.3.2. Extracting additional information

In some cases, non-dominance is a very powerful concept leading to a unique non-dominated alternative. However, in most cases the non-dominated set will be too large to imply a final decision. It may happen that there are several non-dominated alternatives and differences in expected utilities are non-negligible. If such is the case, we should look for additional information that would help us to reduce the classes, and, perhaps, reduce the non-dominated set. Some tools based on functional derivatives to elicit additional information may be seen in Ríos Insua and Ruggeri (2000). Tools based on distance analysis may be seen in Ríos Insua (1990).

4.3.3. Robust solutions

When no additional information may be extracted from experts to reduce the set of inputs, and the set of alternatives is still too big, robust solution concepts should be looked for. One is the maximin approach which may be considered as an automated method which allows to choose actions which guard against catastrophic consequences. A maximin approach would be suitable after a sensitivity analysis has been unable to significantly narrow the range of variation, under changes in inputs of the quantity of interest.

4.3.4. Hyperpriors

Another approach to dealing with lack of robustness would be to place a hyperprior on the class of inputs. Indeed, if there were no possibility of obtaining additional information to deal with the lack of robustness, this technique would be recommended, with the hyperprior being chosen in some default fashion.

5.   Misconceptions in sensitivity analysis

Some basic issues corresponding to a number of sensitivity analysis approaches are described. Their relevance stems from them corresponding to typical misconceptions.

- It is not enough to study changes in output by trying some other inputs.

- Partial sensitivity studies may not be sufficient: a problem may be insensitive to changes in utility and changes in probability, but sensitive to simultaneous changes in utility and probability.

- When performing sensitivity analysis, there are cases in which the optimal value may change a lot, with virtually no change in the optimal action, even if the utility is fixed.

- Alternatively, there are cases in which the optimal alternative varies widely, but the optimal value does not practically change.

- Big changes in optimal value do not necessarily correspond to big changes in consequences of interest.

- Standard global robustness studies, based e.g. on ranges of expected utilities of actions, may not be sufficient within a decision theoretic perspective.

6.   Final comment

Imprecise probability is a generic term used to describe mathematical models that measure uncertainty without precise probabilities. This is certainly the case with robust Bayesian analysis, but there are many other imprecise probability theories, including upper and lower probabilities, belief functions, Choquet capacities, fuzzy logic, and upper and lower previsions; see Walley (1991). Some of these theories, such as fuzzy logic and belief functions, are only tangentially related to sensitivity analysis. Others are intimately related. For example, some classes of probability distributions that are considered in Bayesian sensitivity analysis, such as distribution band classes, can also be interpreted in terms of upper and lower probabilities, see Ríos Insua et al (2000). Also, classes of probability distributions used in sensitivity analysis robust Bayesian analysis will typically generate upper and lower previsions as their upper and lower envelopes, see Berger et al (1996).

Finally, sensitivity analysis is linked to *uncertainty analysis* which aims at quantifying the uncertainty of the output as a function of the uncertainty in the model inputs.

Cross-References

Bayes rule, Bayesian Decision Theory, Decision Analysis, Decision Maker, Linear programming, Multi-criteria decision making, Pareto-optimal solution.

References

Berger J., Betrò B., Moreno E., Pericchi L., Ruggeri F., Salinetti G. and Wasserman L. (eds.) (1996). *Bayesian Robustness*. IMS Lecture Notes.

French S. and Ríos Insua D. (2000). *Statistical Decision Theory*. Arnold.

Gal T. and Greenberg H. (1997). *Advances in Sensitivity Analysis and Parametric Programming*. Kluwer.

Kadane J. and Srinivasan C. (1996). Bayesian Robustness and Stability, in *Bayesian Robustness*, IMS Lecture Notes. Vol. 29, pp. 81-96.

Ríos Insua D. (1990). *Sensitivity Analysis in Multiobjective Decision Making*. Springer.

Ríos Insua D., Ruggeri F. and Martin J. (2000). Bayesian Sensitivity Analysis: A Review. in *Handbook on Sensitivity Analysis* (A. Saltelli et al., eds.). Wiley.

Ríos Insua D. and Ruggeri F. (2000). *Robust Bayesian Analysis*. Springer.

Ruggeri F. and Wasserman L. (1993). Infinitesimal Sensitivity of Posterior Distributions. *Canadian J. Statist.*, 21, pp. 195-203.

Saltelli A., Chan K. and Scott M. (2000). *Mathematical and Statistical Methods: Sensitivity Analysis*. Wiley.

Saltelli A., Tarantola S., Campolongo F. and Ratto M. (2004). *Sensitivity Analysis in Practice: A Guide to Assessing Scientific Models*. Wiley.

Walley P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall.